Testimony of Dario Gil
Vice President AI and IBM Q
Before the U.S. Senate Committee on Science, Commerce, and Transportation
Subcommittee on Communications, Technology, Innovation and the Internet
Russell Senate Office Building – Room 253
Tuesday, December 12, 2017

Introduction

Chairman Wicker, Ranking Member Schatz, members of the subcommittee. Thank you for inviting me here today. My name is Dario Gil and I am Vice President, AI and quantum computing at IBM.

We have arrived at a remarkable moment in the history of information technology. An explosion of data and computation has given us access to massive amounts of digitized knowledge. With it, we have enormous intelligence and power to see patterns and solve problems we never could have before.

Increasingly, the engine we use to tap into this knowledge is artificial intelligence. We can now train algorithms with emerging Artificial Intelligence (AI) technologies to learn directly from data by example. Moreover, we can do this at scale cost through cloud networks to create machines that help humans think. For these reasons, AI is the most important technology in the world today.

The rise of AI has been accompanied by both boundless enthusiasm about its ability to transform our lives, and fears it could potentially harm or displace us. At IBM, we take a different approach. We are guided by the use of artificial intelligence to augment human intelligence. We focus on building practical AI applications that assist people with well-defined tasks. We believe people working collaboratively with these learning systems is the future of expertise.

In my testimony, I'll provide an overview of AI and describe how work in this field has evolved. Then I'll offer some examples of the rapidly growing commercial applications of IBM Watson, the best-known artificial intelligence platform for enterprise business today. I'll also look at how we're beginning to combine AI with other emerging technologies, such as blockchain, to optimize business and provide trust in transactions. I'll describe how AI will impact the nature of work leading to many new and improved job opportunities. Finally, I'll examine IBM's position on the responsible and ethical use of AI.

The Evolution of AI

The idea of creating a ‘thinking’ machine is not new and precedes modern computing. The study of formal reasoning dates to ancient philosophers such as Aristotle and Euclid. Calculating machines were built in antiquity and were improved throughout history by many mathematicians. In the 17th century Leibniz, Hobbes and Descartes explored the possibility that all rational thought could be made as systematic as algebra or geometry.

In 1950, Alan Turing, in his seminal paper *Computing Machinery and Intelligence,* laid out several criteria to assess whether a machine could be deemed intelligent. They have since become known as the “Turing test.” The term "artificial intelligence" was first introduced in 1956, sixty-one years ago, and AI as an academic discipline took off. Three years later, in 1959, IBM scientist Arthur Samuel coined the term “machine learning” to refer to computer algorithms that learn from and make predictions on data by building a model from sample inputs, without following a set of static instructions.

An algorithm is simply a set of rules to be followed in calculations or other problem-solving operations. It can be as basic as the steps involved in solving an addition problem or as complex as instructing a computer how to perform a specific task. One type of machine learning and AI algorithm that has gained tremendous attention over the past several years is an artificial neural network. It has been essential to the explosive growth of AI systems today.

Artificial neural networks are inspired by the architecture of the human brain. They contain many interconnected processing units, called artificial neurons, which are analogous to biological neurons in the brain. Typically, neurons are organized in layers. Different layers may perform different kinds of operations on their inputs. When presented with sample data, an artificial neural network can be trained to perform a specific task, such as recognize speech or images. For example, an algorithm can learn to identify images that contain cars by analyzing numerous images that have been manually labeled as "car" or "no car." It can then use those results to identify cars in images that it has not seen before.

Even though neural networks and other machine learning algorithms were actively researched more than six decades ago, their practical use was hindered by the lack of digitized data from which to learn from and insufficient computational power. At the time, most data were in analog form and not easily available to a computer. Training of the neural network algorithm was and remains a computationally intensive process. Due to the limitations of processors, it could not be implemented effectively.

Over the last decade, the explosion of digital data, the growth in processing speed and power, and the availability of specialized processing devices such as graphical processing units (GPUs) have made it possible to use artificial neural networks in real-world solutions. Today, computation is carried out not only in the cloud and in data centers, but also at the edge of the network, in sensors, wearable devices, smart phones, embedded electronics, factory machines, home devices, or components in a vehicle.

These conditions have also allowed researchers and engineers to create incredibly complex neural networks, called *deep learning* networks. They perform in ways comparable to humans in many tasks. For certain tasks, such as speech and image recognition, game playing, and medical image classification, these networks can outperform people. Today, neural networks are used in a variety of applications, including computer vision, speech recognition, machine translation, social network analysis, playing board and video games, home assistants, conversational devices and chatbots, medical diagnostics, self-driving cars, and operating robots.

In addition to machine learning, AI systems deploy a variety of other algorithms and technologies that include knowledge representation, machine reasoning, planning and scheduling, machine perception (speech and vision), and natural language processing and understanding. At IBM, we are actively researching and advancing these and other technologies so that we can continue to enhance AI systems.

We are also envisioning and developing the next-generation infrastructure required for increasingly complex AI tasks and workloads. This is the physical hardware required to run AI algorithms: the processors, servers, databases, storage, data centers and cloud infrastructure. When all these pieces are aligned in a way that allows algorithms to analyze data with maximum efficiency, we refer to it as the “full stack.”

By successfully engineering the full stack, we can build AI-powered solutions that we can apply to a broad array of societal and industry challenges. While many tend to focus on the benefit of automation, we believe that AI’s true impact will be felt in assisting people’s daily lives, and by helping us carry out extremely complex tasks we cannot do on our own. That includes everything from forecasting the weather, to predicting how traffic will flow, to understanding where crops will grow the best. AI will also help us research the best combinations of compounds for drug development, repurpose chemical structures for new medicines, and optimize vastly intricate supply chains. I’d like to illustrate this further with a look at how IBM Watson is already being used across a range of different industries.

AI applications to industries

My first example illustrates how AI can assist humans in reacting to a problem when there is very little time to react. The security of data and on-line transactions is fundamental to the growth of commerce. But the simple fact is that most organizations can't keep up with the threats. Security experts are too few and overstretched. Sophisticated attacks, including those using AI tools, are coming at a rate that makes them extremely difficult to stop. Entire networks are compromised in the blink of an eye. Watson for cybersecurity allows us to turn the tables. It sifts through the insights contained within vast amounts of unstructured data, whether it's documented software vulnerabilities or the more than 70,000 security research papers and blogs published each year. It instantly alerts security experts to relevant information, scaling and magnifying human cognition. It also learns from each interaction that has an alert, and works proactively to stop continued intrusion. Security analysts, armed with this collective knowledge, can respond to threats with greater confidence and speed.

The second example shows how AI is enhancing customer experience. Tax preparation is an area ripe for AI solutions. H&R Block is using Watson to understand context, interpret intent and draw connections between clients' statements and relevant areas of their tax return. Watson is working alongside H&R Block Tax Pros as they take clients through the tax return process, suggesting key areas where they may qualify for deductions and credits. Clients can follow along and understand how their taxes are being computed and impacted by numerous aspects of the tax code. They can also see the many paths to file a return with the IRS, pinpointing the route that delivers a maximum refund.

A third example demonstrates AI's ability to personalize the client experience. 1-800-Flowers launched an AI-powered gift concierge powered by Watson Conversation. It interacts with online customers using natural language. The service can interpret questions, then ask qualifying questions about the occasion, sentiment and who the gift is for to ensure that suggestions are appropriate and tailored to each customer. In this way, the customer can get the right flower for the right occasion.

The next example highlights AI's role in enhancing agricultural productivity. A program led by our Research division called IBM PAIRS is bringing the power of AI to improve crop yields. It works by processing and analyzing vast amounts of geospatial data to generate vastly improved long term weather and irrigation forecasts. Using these methods, IBM Research and Gallo Winery co-developed a precision irrigation method

and prototype system that allows Gallo to use 20 percent less water for each pound of grapes it produces.

A final example shows AI's ability to optimize supply chains. Traditional brick and mortar retailers are under tremendous pressure from e-commerce. They must find new, cost-effective and efficient ways to deliver goods to buyers in order to stay in business. That means offering customers a range of delivery options – pick up in store, ship from the nearest store, or move goods seamlessly between store and e-commerce. Our client, a major American retailer, had to coordinate this effort across a thousand stores in their fulfillment chain. Our predictive models enabled them to determine optimal distribution across their entire chain, factoring in dozens of different variables. Over an eight-day period including Black Friday and Cyber Monday, they processed more than 4 million orders – a company record -- at a savings of 19% per order compared to the prior year. This led to an overall savings of $7.5 million dollars.

From cybersecurity, to customer experience, to personalization, to productivity and optimization, AI is playing a bigger and bigger role in all realms of commerce. And its uses will only grow.

AI and blockchain

The potential for AI becomes even greater when combined with other emerging technologies such as blockchain. Blockchain stores data in shared, secure, distributed ledgers that allow every participant appropriate access to the entire history of a transaction using a "permissioned" network—one that is highly secure and can determine who can see what. Blockchain holds the promise to be the way we may do all transactions in the future.

A typical AI process could include data collection, use of the data to train a neural network and then deployment of the pre-trained model to power the application. Blockchain supports the AI process by reducing the risk of data tampering and provides data in a form that can be used and audited. There's an old saying in the computer industry "garbage in, garbage out," and that applies to data and how you use it. The integrity of the data used as input to the AI model is a necessary criterion in ensuring the value and usefulness of the model.

Because it can process and analyze massive quantities of data, AI can use blockchain data to gain valuable insights and detect patterns in how supply chains work and processes behave. Over time, this will generate a valuable source of clean, trusted transactional data that cuts across industries to give us new insights. That includes both structured and unstructured data -- everything from Internet of Things (IoT)

information to compliance and geospatial data that's stored on a blockchain. AI can use this information to generate valuable insights and detect patterns in near-real time, driving new efficiencies across business operations.

For example, IBM Research is working with Everledger, a company that tracks and protects diamonds and other valuables. We're using AI to analyze digital information on one million diamonds Everledger has stored on a blockchain. We can cross-check that data against UN regulations to prevent the sale of conflict diamonds. We can verify time and date stamps. We can certify laser inscriptions in the girdle of the stone. We can perform these analytics directly on the blockchain, without the need to extract the data first. This minimizes opportunities for data tampering and fraud. While this is a specialized application, it shows some of the kinds of data we can collect and analyze at huge scale.

We have also partnered with Walmart to use blockchain and AI techniques to ensure food safety. Today's food supply chains are highly complex and involve multiple components, stakeholders, and activities. This complexity makes it difficult to identify sources of contamination, counterfeit substitutions, loss of refrigeration, or food transportation safety issues as products move from their sources to their consumption by consumers. Blockchain supports traceability by tracking the food products from origin to destination and by allowing certification of respective transactions and events along the way. AI-powered technologies are used to analyze this information to help ensure food that can be eaten safely.

AI and the Future of Work

Artificial intelligence will alter the way people work. This has been true of many new technologies that have benefited human populations over time because they dramatically improved industrial output. They have led to fewer grueling jobs. In the process, new types of jobs have emerged. However, such disruptive improvements have always called for a period of training and adjustment.

We need to openly understand and recognize this fact, so that we can create the right conditions to make this transition as successful as possible. As a nation, we need to be prepared to offer the appropriate education and support to manage this change well. There's no question the advent of artificial intelligence will impact jobs. Despite the fear, anxiety, and prediction of massive job loss, history suggests that, even in the face of technological transformation, employment continues to grow and very few occupations disappear.

Rather, it is the transformation of occupations that is very likely to be widespread that will impact most workers. Occupations are made up of tasks. It is the tasks that are automated and reorganized where the transformation occurs. Workers will need new skills for the new transformed tasks and occupations. But, it is the tasks that cannot or will not be automated where workers provide the greatest value, commanding higher wages and incomes as a result.

Some “new collar jobs” will emerge - jobs that require advanced technical skills but do not necessarily require a full undergraduate education. A study by Accenture of more than 1,000 large companies that are already using or testing AI and machine-learning systems identified the emergence of entire categories of new, uniquely human jobs with no precedents.

For example, “trainers” will be required to teach AI systems how they should perform. They may write scripts for chatbots, helping them to improve their understanding of human communication, or help provide labeled data needed to train the algorithm. They may teach AI systems how to personalize their recommendations, or show compassion when responding to certain problems that require it. Another category of “explainers” will be needed to help convey how AI systems have arrived at a decision or a recommendation. They’ll monitor the operations of AI systems or perform forensic analyses on algorithms and make corrections to them if they generate an incorrect result. Earlier, I referenced the shortage of qualified cybersecurity professionals. In the future, we’ll need far more of them to engage with AI systems, review the recommendations these systems offer and act decisively upon threats.

There are actions we must take now to ensure the workforce is prepared to embrace the era of AI and the ways it will augment our economy. To begin, we must address the shortage of workers with the skills needed to make advances in AI, create new solutions and work in partnership with AI systems. We need to match skills education and training with the actual skills that will be required in the emerging age of AI.

At IBM, we have an educational model called P-TECH to train new collar workers for a job in technology. P-TECH combines the best of high-school, community college, hands-on-skills training and professional mentoring, and provides public high school students in grades 9-14 a path to post-graduation opportunities in fields aligned with the skills American employers are looking for.

Our goal must be to create multiple pathways like this for more people to acquire the skills that will be in demand, as AI use becomes more commonplace. We can use the example of the adoption of software programming as a critical skill that is taught in many high school and colleges. Some colleges require that all students learn how to

code since they consider it a necessary skill for success. Students becoming proficient in programming have a wider range of job opportunities.

In the future, we may promote and see a similar trend with students gaining understanding of and proficiency in AI techniques such as machine learning. Preparing more U.S. students and workers for success in these well-paying new collar jobs is essential if we want a workforce that is ready to capitalize fully on AI's economic promise.

Let me also say that as well-intentioned as it may seem to some, taxing automation will not serve the cause of fostering employment in the new AI economy. It will only penalize technological progress. We should not adopt measures like this one that will harm America's competitiveness.

Inevitably, people adapt best by finding higher value in new skills. Technologies that are easiest to integrate and integrate with will be those that improve human productivity. But they should not replace human judgment. IBM Watson was designed from the beginning to work in concert with human expertise. It will only be successful as long as there are people with the right skills to engage with it.

Building trust in AI

To enjoy the full benefits of AI, we will also need to have confidence in the recommendations, judgments and uses of AI systems. IBM is deeply committed to the responsible and ethical development of AI. Last year, we published one of the first corporate white papers on this subject. The paper, which was intended to help launch a global conversation, centered around the need for safe, ethical, and socially beneficial management of AI systems.

Trust in automated systems is not a new concept. We drive cars trusting the brakes will work when the pedal is pressed. We perform laser eye surgery trusting the system to make the right decisions. We have automated systems fly our airplanes trusting they will navigate the air correctly. In these cases, trust comes from confidence that the system will not make a mistake, leveraging system training, exhaustive testing, and experience. We will require similar levels of trust for AI systems, applying these methodologies.

In some cases, users of AI systems will need to justify why an AI system produced its recommendation. For example, doctors and clinicians using AI systems to support medical decision-making may be required to provide specific explanations for a diagnosis or course of treatment, both for regulatory and liability reasons. Thus, in

these cases, the system will need to provide the reasoning and motivations behind the recommendation, in line with existing regulatory requirements specific to that industry. In the European Union, this will be a requirement for all automated decision-making AI systems as of May 2018.

These safeguards can also help to manage the potential for bias in the decision-making process, another important concern with AI. Bias can be introduced both in the datasets that are used to train an AI system and by the algorithms that process that data, and how people interpret and communicate the discerned insights. Our belief is that the data and algorithmic aspects can not only be managed, but also that AI systems themselves can help eliminate many of the biases that already exist in human decision-making models today.

At the beginning of this year, IBM issued principles for transparency and trust to guide our development and use of AI systems. In summary, they state the following:

- We believe AI's purpose is to augment human intelligence
- We will be transparent about when and where AI is being applied, and about the data and training that went into its recommendations.
- We believe our clients' data and insights are theirs.
- We are committed to helping students, workers, and citizens acquire the skills to engage safely, securely, and effectively with cognitive systems, and to do the new kinds of work that will emerge in an AI economy.

In the same way that we are committed to the responsible use of AI systems, we are committed to the responsible stewardship of the data they collect. We also believe that government data policies should be fair and equitable and prioritize openness.

IBM is actively innovating in this field. We are deriving best practices for how, when, and where AI algorithms should be used. We are creating new AI algorithms that are more explainable and more accurate. We are working on the algorithmic underpinnings of bias and AI, such as creating technologies that can identify and cleanse illegal biases from training datasets.

Of course, no single company can guarantee the safe and responsible use of such a pervasive technology. For this reason, IBM is a founding member of the Partnership on AI, a collaboration with other key industry leaders and many scientific and nonprofit organizations. Its goal is to share best practices on AI technologies, advance the public's understanding, and serve as an open platform for discussion and engagement about AI and its influences on people and society.

AI has enormous transformative power. Much has been said about its potential to transform sectors and industries. However, AI is also giving us a technological toolkit to address many societal challenges. At IBM we are committed to pioneering new solutions, and showcasing and promoting the opportunities to use AI in social good applications. Three years ago, we launched the AI for Social Good program and have executed a number of AI for Good projects, from using AI to understand patterns of opioid addiction, to prototyping recommendation systems that would aid low-income individuals and help them stay out of poverty, to applying machine learning to understand transmission mechanisms of the Zika virus.

Earlier this year, we announced the MIT-IBM Watson AI Lab, a partnership with Massachusetts Institute of Technology (MIT) to carry out fundamental AI research. One of the research areas for the lab is focused on *advancing shared prosperity through AI* - exploring how AI can deliver economic and societal benefits to a broader range of people, nations and enterprises.

Lastly, no discussion of the future of AI would be complete without acknowledging the critical role of government. Public investment and policy support have been the twin pillars of American global technological leadership for the past half-century. We hope and expect the same will be true in the coming age of AI. For this reason, we enthusiastically welcome the interest and support of the United States Senate as this technology continues to evolve. Together, we can ensure that AI serves people at every level of society and advances the common good.